arXiv:1206.0992v2 [cs.DC] 18 Jun 2012

When Do Gossip Algorithms Converge in Finite Time?[*]

Guodong Shi, Mikael Johansson and Karl Henrik Johansson[†]

Abstract

In this paper, we show that there exists a symmetric gossip algorithm that converges in finite time if and only if the number of network nodes is a power of two. We also show that there always exists a globally finite-time convergent gossip algorithm despite the number of nodes if asymmetric gossiping is allowed. In addition, an "all-or-nothing" theorem is established on the finite-time convergence of distributed averaging algorithms: for a class of averaging algorithms defined by matrix sequences selected from a countable set of stochastic matrices, finite-time convergence either holds for all initial conditions, or fails for almost all initial conditions.

Keywords: Gossip algorithms, Finite-time convergence, Distributed averaging

1 Introduction

Various gossip algorithms, in which information exchange is always carried out pairwise among the nodes, have been widely used to structure distributed computation, optimization, and signal processing over peer-to-peer, sensor, and social networks [3, 2, 8, 5, 10, 11, 12, 13, 6, 7]. Gossip averaging plays a fundamental role in the study of gossip algorithms due to its simple nature and wide application.

Consider a network with node set $V = \{1, \dots, n\}$. Let the value of node i at time k be

---

[*] This work has been supported in part by the Knut and Alice Wallenberg Foundation, the Swedish Research Council and KTH SRA TNG.

[†] The authors are with ACCESS Linnaeus Centre, School of Electrical Engineering, Royal Institute of Technology, Stockholm 10044, Sweden. Email: {`guodongs, mikaelj, kallej`}`@kth.se`

$x_i(k) \in \mathbb{R}^1$ for $k \geq 0$. Let $P_0, \dots, P_k, \dots$ be a sequence of matrices satisfying

$$P_k \in \mathcal{M} \doteq \Big\{ M_{ij} \doteq I - \frac{(e_i - e_j)(e_i - e_j)^T}{2} \ : \ i, j = 1, \dots, n \Big\},$$

where $e_m = (0 \dots 0 \ 1 \ 0 \dots 0)^T$ is the $n \times 1$ unit vector whose m'th component is 1. Denote $x(k) = (x_1(k) \dots x_n(k))^T$. Then a symmetric deterministic gossip algorithm is defined by

$$x(k+1) = P_k x(k). \tag{1}$$

Algorithm (1) and its variations have been extensively studied in the literature for both randomized and deterministic models. Karp et al. [2] derived a general lower bound for synchronous gossiping; Kempe et al. [3] proposed a randomized gossiping algorithm on complete graphs and determined the order of its convergence rate. Then Boyd et al. [5] established both lower and upper bounds for the convergence time of synchronous and asynchronous randomized gossiping algorithms, and developed algorithms for optimizing parameters to obtain fast consensus. Fagnani and Zampieri discussed asymmetric gossiping in [8]. Liu et al. [9] presented a comprehensive analysis for the asymptotic convergence rates of deterministic averaging, and recently distributed gossip averaging subject to quantization constraints was studied in [12]. Distributed signal processing and estimation algorithms via gossiping were discussed in [10, 11]. A detailed introduction to gossip algorithms can be found in [6]. Few works have considered the possibility of finite-time convergence of gossip algorithms.

We introduce the following definition for finite-time convergence.

Definition 1 *A gossip algorithm achieves finite-time convergence with respect to initial value* $x(0) = x^0 \in \mathbb{R}^n$ *if there exists an integer* $T(x^0) \geq 0$ *such that* $x(T) \in \mathrm{span}\{\mathbf{1}\}$*. Global finite-time convergence is achieved if such* $T(x^0)$ *exists for every initial value* $x^0 \in \mathbb{R}^n$*.*

The aim of this paper is to establish the fundamental limit of gossip averaging in terms of finite-time convergence. Consider the following simple examples.

Example 1 *Consider a network with four nodes* $V = \{1, \dots, 4\}$ *and initial values* $x_1(0), \dots, x_4(0)$*. We define a gossip algorithm* $\{P_k\}$ *by*

$$P_0 = M_{12}, \ P_1 = M_{34}, \ P_3 = M_{13}, \ P_4 = M_{24}$$

following the node pair order $(1,2)$*,* $(3,4)$*,* $(1,3)$*,* $(2,4)$*, as illustrated in Figure 1. Then all the nodes reach the same value* $\sum_{i=1}^{4} x_i(0)/4$ *in four steps. Therefore, finite-time convergence is achieved.*

Example 2 *Consider instead a network with three nodes* $V = \{1, \dots, 3\}$ *and initial values* $x_1(0), \dots, x_3(0)$*. Then if consensus cannot be reached in one step, a common value can never be reached in finite time via gossip averaging because there will always be two different values in the network, see Figure 2. Note that the initial values from which consensus can be reached in one step are contained in the following set:*

$$\big\{2x_1(0) = x_2(0) + x_3(0)\big\} \bigcup \big\{2x_2(0) = x_1(0) + x_3(0)\big\} \bigcup \big\{2x_3(0) = x_1(0) + x_2(0)\big\},$$

which is of measure zero under standard Lebesgue measure in $\mathbb{R}^3$*. Therefore, finite-time convergence by gossiping is impossible for three nodes for almost all initial values.*

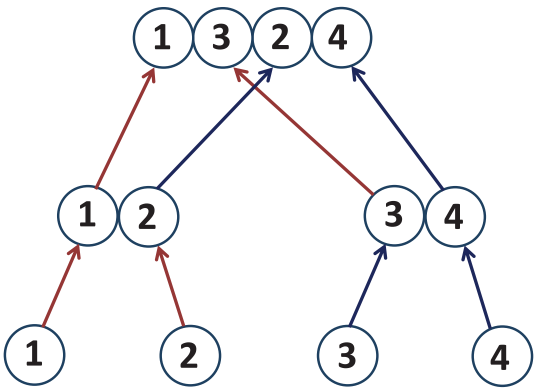

Figure 1: Consensus can always be achieved in four steps for all initial values by gossiping with four nodes.

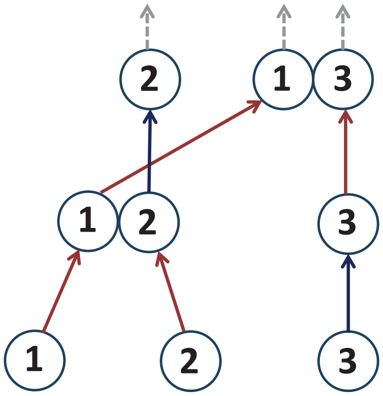

Figure 2: Consensus can never be achieved in finite time for almost all initial values by gossiping with three nodes.

We see from the examples that the possibility of finite-time convergence via gossiping indeed relies on the number of nodes. We present the following main result on the finite-time convergence of gossip algorithms.

Theorem 1 *There exists a deterministic gossip algorithm* $\{P_k\}_0^\infty$*,* $P_k \in \mathcal{M}, k \geq 0$*, that converges globally in finite time if and only if there exists an integer* $m \geq 0$ *such that* $n = 2^m$*.*

In fact, as will be shown in the proof of Theorem 1, if $n = 2^m$ for some integer $m \geq 0$, we can construct a gossip algorithm which converges in $(n \log n)/2$ steps. The idea for constructing such an algorithm follows from the node pair selection strategy in Example 1.

Recall that a finite square matrix $M = [m_{ij}] \in \mathbb{R}^{n \times n}$ is called *stochastic* if $m_{ij} \geq 0$ for all i, j and $\sum_j m_{ij} = 1$ for all i [1]. Indeed, (1) is a special case of distributed averaging algorithms defined by products of stochastic matrices [15, 16, 17]:

$$x(k+1) = W_k x(k), \tag{2}$$

where $W_k \in \mathcal{S} \doteq \big\{W \in \mathbb{R}^{n\times n}:\ W \text{ is a stochastic matrix}\big\}$. Let $\mathcal{S}_0 \subseteq \mathcal{S}$ be a subset of stochastic matrices. We define

$$\mathscr{X}_{\mathcal{S}_0} \doteq \Big\{x \in \mathbb{R}^n:\ \exists W_0,\dots,W_s \in \mathcal{S}_0, s \geq 0\ \text{ s.t. } W_s\cdots W_0 x \in \operatorname{span}\{\mathbf{1}\}\Big\}.$$

Let $\mathbf{M}(\cdot)$ represent the standard Lebesgue measure on $\mathbb{R}^n$. We have the following result for the finite-time convergence of general averaging algorithms.

Theorem 2 *(All or Nothing) Suppose $\mathcal{S}_0$ is a set with at most countable elements. Then either $\mathscr{X}_{\mathcal{S}_0} = \mathbb{R}^n$ or $\mathbf{M}(\mathscr{X}_{\mathcal{S}_0}) = 0$. In fact, if $\mathscr{X}_{\mathcal{S}_0} \neq \mathbb{R}^n$, then $\mathscr{X}_{\mathcal{S}_0}$ is a union of at most countably many linear spaces whose dimensions are no larger than $n-1$.*

Theorem 2 implies, given countably many stochastic matrices contained in a set $\mathcal{S}_0$, either for any initial value $x^0 \in \mathbb{R}^n$, we can select a sequence of matrices from $\mathcal{S}_0$ so that the obtained averaging algorithm converges in finite time starting from x^0, or for almost all initial values, any averaging algorithm obtained by sequence selection from $\mathcal{S}_0$ fails to converge in finite time.

For a fixed averaging algorithm, $\{W_k\}_0^\infty$ with $W_k \in \mathcal{S}, k \geq 0$, taking $\mathcal{S}_0 = \{W_k, k \geq 0\}$, it can be immediately concluded from Theorem 2 that either this algorithm converges in finite time for all initial values, or fails to converge in finite time for almost all initial values.

Noticing that $\mathcal{M}$ is a finite set and utilizing Theorems 1 and 2, we obtain the following conclusion.

Theorem 3 *Suppose there exists no integer $m \geq 0$ such that $n = 2^m$. Then for almost all initial values, it is impossible to find a gossip algorithm $\{P_k\}_0^\infty$ with $P_k \in \mathcal{M}, k \geq 0$, to reach finite-time convergence.*

The relation between Theorem 1 and Theorem 3 is as follows. Theorem 1 indicates that if the number of nodes n is not a power of two, finding a gossip algorithm which converges globally in finite time is impossible. However, in this case, there still might exist a gossip algorithm which converges in finite time for some initial values, say, half of $\mathbb{R}^n$. Theorem 3 thus further excludes the possibility of the existence of such algorithms by an indeed stronger claim, which shows that the initial values from which there exists a gossip algorithm converging in finite time form a measure zero set.

We give some remarks on randomized algorithms. Most existing works on gossiping algorithms use randomized models [3, 2, 8, 5, 10, 11, 12, 13]. Deterministic gossiping was discussed in [12, 9]. Although we consider deterministic algorithms in this paper, the results can still be easily extended to randomized gossip algorithms. Note that Theorems 1, 2 and 3 do not focus on any particular algorithm, but on the class of algorithms defined on $\mathcal{M}$ or $\mathcal{S}_0$. Therefore, even if some randomized structure is introduced on $\{P_k\}_0^\infty$ and $\{W_k\}_0^\infty$ in (1) and (2), our conclusions still stand in the almost sure sense. Moreover, it is worth pointing out that even if $\mathcal{S}_0$ is finite (as $\mathcal{M}$), there will still be uncountably many different averaging algorithms in the form of (2) as long as $\mathcal{S}_0$ contains at least two elements.

From Theorem 3, we can conclude that finite-time convergence is rare for symmetric gossiping since in practice the number of nodes can hardly be some power to two. Now we relax the requirement of gossiping in way that asymmetric updates are also allowed [8]. Introduce

$$\mathcal{M}_* \doteq \Big\{ I - \frac{(e_i - e_j)(e_i - e_j)^T}{2} : \ i, j = 1, \dots, n \Big\} \bigcup \Big\{ I - \frac{e_i (e_i - e_j)^T}{2} : \ i, j = 1, \dots, n \Big\}.$$

Note that matrix $I - e_i(e_i - e_j)^T/2$ represents the case that node i updates only when it is gossiping with node j.

It turns out that finite-time convergence is always possible despite the number of nodes as long as asymmetric gossiping is allowed. The following conclusion holds.

Theorem 4 *For any number of nodes, there always exists a deterministic gossip algorithm* $\{P_k\}_0^\infty$, $P_k \in \mathcal{M}_*, k \geq 0$, *which converges globally in finite time.*

2 Proofs

We prove Theorems 1, 2 and 4 in this section.

2.1 Proof of Theorem 1

(Necessity.) Suppose $n = 2^{n_1} n_2$ with $n_1 \geq 0$ and $n_2 \geq 3$ an odd integer. Suppose $P_0, \dots, P_{k_*} \in \mathcal{M}$ with $k_* \geq 0$ gives an algorithm of (1) that converges in finite time globally.

Take $x_1, \dots, x_{2^{n_1}} = 0$ and $x_{2^{n_1}+1}, \dots, x_n = 2^{k_*+1}$. Then there exists $c \in \mathbb{R}$ such that $x_i(k_*+1) = c, i = 1, \dots, n$. Because each element in $\mathcal{M}$ is symmetric and therefore doubly stochastic, average is always preserved. Thus, we have

$$c = \frac{2^{k_*+1}2^{n_1}(n_2-1)}{2^{n_1}n_2} = \frac{2^{k_*+1}(n_2-1)}{n_2}.$$

On the hand, it is not hard to see that c is an integer for the giving initial value since pairwise averaging takes place k_*+1 times. Consequently, we have $c = r_2 2^{r_1}$ with $0 \leq r_1 \leq k_*+1$ an integer and $r_2 \geq 1$ an odd integer.

Therefore, we conclude that

$$\frac{2^{k_*+1}(n_2-1)}{n_2} = r_2 2^{r_1},$$

which implies

$$2^{k_*+1-r_1}(n_2-1) = r_2 n_2. \tag{3}$$

This is impossible because the left-hand side of Eq. (3) is an even number while the right-hand side odd. Therefore, (1) cannot achieve global finite-time convergence no matter how $P_0, \dots, P_k, \dots$ are chosen.

(Sufficiency.) We need to construct a gossip algorithm which converges in finite time globally for $n = 2^m$.

We relabel the nodes in a binary system. We use the binary number

$$B_1 \dots B_m, \quad B_s \in \{0, 1\}, s = 1, \dots, m$$

to mark node i if $B_1 \dots B_m = i - 1$ as a binary number. The gossip algorithm is derived from the following matrix selection process:

S1. Let $k = 1$.

S2. Take 2^{m-1} matrices from $\mathcal{M}$, as the elements in the following set

$$\mathcal{P}_k \doteq \Big\{ I - \frac{(e_i - e_j)(e_i - e_j)^T}{2} : \ \text{in the binary system, the } k\text{'th digit of } i-1 \text{ equals } 0, \text{ and the } k\text{'th digit of } j-1 \text{ equals } 1 \Big\}.$$

In other words, we take all the node pairs (i, j), where $i-1$ and $j-1$ have identical expressions in the binary system except for the k'th digit. Label the matrices in $\mathcal{P}_k$ as $P^*_{(k-1)2^{m-1}}, \dots, P^*_{k2^{m-1}-1}$ with an arbitrary order.

S3. Let $k = k + 1$ and go to S2 until $k = m$.

Following this matrix selection process, $P_0^*, \dots, P_{m2^{m-1}-1}^*$ gives a gossip algorithm in the form of (1). It is easy to see that the vector

$$P_{s2^{m-1}-1}^* \cdots P_0^* x^0, \quad x^0 \in \mathbb{R}^n, \ s = 1, \dots, m$$

has at most 2^{m-s} different elements. Thus, convergence is reached after $m2^{m-1} = (n \log_2 n)/2$ updates. This completes the proof.

Remark 1 *We conjecture that for $n = 2^m$ nodes, the gossip algorithm defined by the matrix selection process given in the proof actually reaches the fastest convergence. In other words, global finite-time convergence cannot be achieved in less than $(n \log_2 n)/2$ steps.*

2.2 Proof of Theorem 2

Define a function $\delta(M)$ of a matrix $M = [m_{ij}] \in \mathbb{R}^{n\times n}$ by (cf. [14])

$$\delta(M) \doteq \max_j \max_{\alpha,\beta} |m_{\alpha j} - m_{\beta j}|. \tag{4}$$

Given an averaging algorithm (2) defined by $\{W_k\}_0^\infty$ with $W_k \in \mathcal{S}_0, k \geq 0$. Suppose there exists an initial value $x^0 \in \mathbb{R}^n$ for which $\{W_k\}_0^\infty$ fails to achieve finite-time convergence. Then obviously $\delta(W_s \cdots W_0) > 0$ for all $s \geq 0$.

Claim. $\mathrm{rank}(W_s \cdots W_0) \geq 2, \ s \geq 0$.

Let $W_s \cdots W_0 = (\omega_1 \dots \omega_n)^T$ with $\omega_i \in \mathbb{R}^n$. Since $\delta(W_s \cdots W_0) > 0$, there must be two rows in $W_s \cdots W_0$ that are not equal. Say, $\omega_1 \neq \omega_2$. Note that $W_s \cdots W_0$ is a stochastic matrix because any product of stochastic matrices is still a stochastic matrix. Thus, $\omega_i \neq 0$ for all $i = 1, \dots, n$. On the other hand, if $\omega_1 = c\omega_2$ for some scalar c, we have $1 = \omega_1^T \mathbf{1} = c\omega_2^T \mathbf{1} = c$, which is impossible because $\omega_1 \neq \omega_2$. Therefore, we conclude that $\mathrm{rank}(W_s \cdots W_0) \geq \mathrm{rank}(\mathrm{span}\{\omega_1, \omega_2\}) \geq 2$. The claim holds.

Suppose there exists some $y \in \mathbb{R}^n$ such that $y \notin \mathscr{X}_{\mathcal{S}_0}$. We see from the claim that the dimension of $\ker(W_s \cdots W_0)$ is at most $n - 2$ for all $s \geq 0$ and $W_0, \dots, W_s \in \mathcal{S}_0$.

Now for $s = 0, 1, \dots,$ introduce

$$\Theta_s \doteq \big\{ x \in \mathbb{R}^n : \ \exists W_0, \dots, W_s \in \mathcal{S}_0, \ \text{s.t. } W_s \cdots W_0 x \in \mathrm{span}\{\mathbf{1}\} \big\}.$$

Then Θ_s indicates the initial values from which convergence is reached in $s+1$ steps. For any fixed $W_0, \dots, W_s \in \mathcal{S}_0$, we define

$$\Upsilon_{W_s \dots W_0} \doteq \big\{z \in \mathbb{R}^n : \quad W_s \cdots W_0 z \in \text{span}\{\mathbf{1}\}\big\}.$$

Clearly $\Upsilon_{W_s \dots W_0}$ is a linear space. It is straightforward to see that $\Theta_s = \bigcup_{W_s \dots W_0 \in \mathcal{S}_0} \Upsilon_{W_s \dots W_0}$, and therefore

$$\mathscr{X}_{\mathcal{S}_0} = \bigcup_{s=0}^{\infty} \Theta_s = \bigcup_{s=0}^{\infty} \bigcup_{W_s, \dots, W_0 \in \mathcal{S}_0} \Upsilon_{W_s \dots W_0}.$$

Noticing that $z \in \Upsilon_{W_s \dots W_0}$ implies $\big(z - W_s \cdots W_0 z\big) \in \ker(W_s \cdots W_0)$, we define a linear mapping

$$\begin{aligned} f: \quad \Upsilon_{W_s \dots W_0} &\longmapsto \ker(W_s \cdots W_0) \times \text{span}\{\mathbf{1}\} \\ \text{s.t.} \quad f(z) &= \big(z - W_s \cdots W_0 z, W_s \cdots W_0 z\big) \end{aligned} \tag{5}$$

Suppose $z_1, z_2 \in \Upsilon_{W_s \dots W_0}$ with $z_1 \neq z_2$. It is straightforward to see that either $W_s \cdots W_0 z_1 = W_s \cdots W_0 z_2$ or $W_s \cdots W_0 z_1 \neq W_s \cdots W_0 z_2$ implies $f(z_1) \neq f(z_2)$. Hence, f is injective. Therefore, noting that $\ker(W_s \cdots W_0)$ is a linear space with dimension at most $n-2$, we have $\dim(\Upsilon_{W_s \dots W_0}) \leq n-1$, and thus $\mathbf{M}(\Upsilon_{W_s \dots W_0}) = 0$. Consequently, we conclude that

$$\mathbf{M}(\Theta_s) = \mathbf{M}\Big(\bigcup_{W_0, \dots, W_s \in \mathcal{S}_0} \Upsilon_{W_s \dots W_0}\Big) \leq \sum_{W_0, \dots, W_s \in \mathcal{S}_0} \mathbf{M}\big(\Upsilon_{W_s \dots W_0}\big) = 0$$

because any finite power set $\mathcal{S}_0 \times \cdots \times \mathcal{S}_0$ is still a countable set as long as $\mathcal{S}_0$ is countable. This immediately leads to

$$\mathbf{M}(\mathscr{X}_{\mathcal{S}_0}) = \mathbf{M}\Big(\bigcup_{s=0}^{\infty} \Theta_s\Big) \leq \sum_{s=0}^{\infty} \mathbf{M}(\Theta_s) = 0.$$

Additionally, since every Θ_s is a union of at most countably many linear spaces, each of dimension no more than $n-1$, $\mathscr{X}_{\mathcal{S}_0}$ is also a union of countably many linear spaces with dimension no more than $n-1$. The desired conclusion thus follows.

2.3 Proof of Theorem 4

We first establish some auxiliary lemmas.

Lemma 1 *Let $n \geq 1$ be an integer. Then there exists a unique expression of n in the form of*

$$n = \sum_{k=0}^{D(n)} \theta_k 2^k, \tag{6}$$

where $D(n) \leq \lfloor (n-1)/2 \rfloor$ is an integer with $\lfloor z \rfloor$ representing the largest integer no lager than z, and $\theta_k \in \{0, 1\}$ for all k.

Proof. We first prove the existence and then the uniqueness.

(Existence.) Take $k_1 = \max\{k \in \mathbb{Z} : \; 2^k \leq n\}$. Then $n_1 = n - 2^{k_1} \geq 0$ is a nonnegative integer. If $n_1 = 0$, the desired expression for n has been found. Otherwise, we continue to take $k_2 = \max\{k \in \mathbb{Z} : \; 2^k \leq n_1\}$ and define $n_2 = n_1 - 2^{k_2} \geq 0$. Again, the proof is ended if $n_2 = 0$ and otherwise the process can be proceeded to define $k_3, k_4, \ldots$. Since n is a finite number, this process cannot last for infinity so that the existence of (6) for n has been proved. Moreover, it is straightforward to see that

$$2(D(n)+1) \leq \sum_{k=0}^{D(n)} \theta_k 2^k + 1 = n + 1,$$

which implies $D(n) \leq \lfloor (n-1)/2 \rfloor$.

(Uniqueness.) Suppose there exist integers $D \geq 1$ and $\theta_k, \hat{\theta}_k \in \{0, 1\}$ such that

$$\sum_{k=0}^{D} \theta_k 2^k = \sum_{k=0}^{D} \hat{\theta}_k 2^k. \tag{7}$$

Respectively, we define

$$\zeta_* \doteq \max\{k : \; \theta_k = 1, k = 0, \ldots, D\}; \quad \hat{\zeta}_* \doteq \max\{k : \; \hat{\theta}_k = 0, k = 1, \ldots, D\}.$$

Claim. We have $\zeta_* = \hat{\zeta}_*$.

We prove this claim by contradiction. Say ζ_* and $\hat{\zeta}_*$ are not equal. Without loss of generality we assume $\zeta_* > \hat{\zeta}_*$. According to (7), this leads to

$$2^{\zeta_*} = \sum_{k=0}^{\zeta_*-1} \hat{\theta}_k 2^k - \sum_{k=0}^{\zeta_*-1} \theta_k 2^k \doteq \sum_{k=0}^{\zeta_*-1} \chi_k 2^k, \tag{8}$$

where $\chi_k \in \{0, 1, -1\}$. Now we take $\varsigma_* = \min\{k : \chi_k \neq 0, k = 0, \ldots, \zeta_* - 1\}$. Then (8) leads to

$$2^{\zeta_*} = \sum_{k=0}^{\zeta_*-1} \chi_k 2^k = 2^{\varsigma_*}\Big(\sum_{k=\varsigma_*}^{\zeta_*-1} \chi_k 2^{k-\varsigma_*}\Big). \tag{9}$$

Noting the fact that

$$\sum_{k=\varsigma_*}^{\zeta_*-1} \chi_k 2^{k-\varsigma_*} = \chi_{\varsigma_*} + \sum_{k=\varsigma_*+1}^{\zeta_*-1} \chi_k 2^{k-\varsigma_*}$$

is an odd number, a contradiction aries that in (9), the left-hand side is even while the right-hand side odd. This proves the claim.

Since $\zeta_* = \hat{\zeta}_*$, we have remove the term 2^{ζ_*} from both the sides of (7) so that

$$\sum_{k=0}^{\zeta_*-1} \theta_k 2^k = \sum_{k=0}^{\zeta_*-1} \hat{\theta}_k 2^k. \tag{10}$$

Repeating the same argument we know that the largest powers of the nonzero items in the right-hand and left-hand sides of (10) are the same. This process continues until we can verify that $\theta_k = \hat{\theta}_k$ for all k in (7). This proves the uniqueness of the expression we construct in the existence proof.

The proof is complete. $\square$

Lemma 2 *For n nodes, there exists a deterministic gossip algorithm $\{P_k\}_0^\infty$, $P_k \in \mathcal{M}_*, k \geq 0$, which converges in finite time for the initial condition $x_2(0) = \cdots = x_n(0)$.*

Proof. Denote $x_1(0) = a$ and $x_2(0) = \cdots = x_n(0) = b$. Construct a gossip algorithm as follows:

(i) For $j = 2, \dots, n-1$, node j gossips with node 1 asymmetrically, i.e., node 1 stays at a, and node j updates to $(a+b)/2$.

(ii) Node 1 and node n do asymmetric gossiping so that they all update to $(a+b)/2$.

We see all the node state converge to $(a+b)/2$ in $n-1$ iterations. The proof is complete. $\square$

We are now in a place to prove the theorem by an inductive approach.

Denote $|S|$ for a set S as the number of elements in S. Based on Lemma 1, we can well define a function as follows:

$$\begin{aligned} \hbar:\quad & \mathbb{Z}_+ \to \mathbb{Z}_+ \\ \text{s.t.}\quad & \hbar(n) = \big|\{\theta_0, \dots, \theta_{D_n}\}\big| \text{ with } n = \textstyle\sum_{k=0}^{D(n)} \theta_k 2^k. \end{aligned} \tag{11}$$

We divide the rest of the proof into three steps.

Step 1. For the number of nodes n satisfying $\hbar(n) = 1$, it has been proved in Theorem 1 that there exists a globally convergent gossip algorithm using symmetric updates.

Step 2. Suppose the theorem holds for all n satisfying $\hbar(n) = K$ with $K \geq 1$ an integer. This is to say, there exists a gossip algorithm, $\{P_k\}_0^\infty$, $P_k \in \mathcal{M}_*, k \geq 0$, converging in finite time globally if $\hbar(n) = K$.

Step 3. Now let $\hbar(n) = K + 1$. In this case, we have

$$n = \sum_{k=1}^{K+1} 2^{d_k}, \ d_k \in \mathbb{Z}_{\geq 0}. \tag{12}$$

Without loss of generality we assume $d_1 < d_2 < \cdots < d_{K+1}$. Then we have

$$n = 2^{d_1}\Big(1 + \sum_{k=2}^{K+1} 2^{d_k - d_1}\Big). \tag{13}$$

Now based on Step 2, for $n_1 = \sum_{k=2}^{K+1} 2^{d_k-d_1}$ nodes, there exists a gossip algorithm converging in finite time globally in light of the fact that $\hbar(n_1) = K$. Lemma 2 further ensures that $1 + \sum_{k=2}^{K+1} 2^{d_k-d_1}$ nodes can converge globally in finite time.

For $n = 2^{d_1}\Big(1 + \sum_{k=2}^{K+1} 2^{d_k-d_1}\Big)$ nodes, we can divide the nodes into $n_1 = 1 + \sum_{k=2}^{K+1} 2^{d_k-d_1}$ groups, $V_1, \dots, V_{n_1}$, each of which has 2^{d_1} nodes. We first make the nodes in the same group reach the same value in finite time, which is possible based on Theorem 1. Next, we reorganize the nodes into 2^{d_1} groups, $\hat{V}_1, \dots, \hat{V}_{2^{d_1}}$, each of which has n_1 nodes picked from $V_1, \dots, V_{n_1}$, respectively. In this way, the initial values for $\hat{V}_1, \dots, \hat{V}_{2^{d_1}}$ are just equivalent copies. Thus, finite-time convergence can be reached when the same algorithm (with respect to node values) is applied to these groups.

Therefore, we have proved that there exists a gossip algorithm $\{P_k\}_0^\infty$, $P_k \in \mathcal{M}_*, k \geq 0$, which converges globally in finite time for all n satisfying $\hbar(n) = K + 1$.

We can finally conclude from the principle of mathematical induction that the desired conclusion holds for all $\hbar(n)$, and thus for all n. This completes the proof.

3 Conclusions

In this paper, we proved that $n = 2^m$ for some integer m is a critical finite-time convergence condition for symmetric gossip averaging algorithm on a network of size n. If there exists some integer $m \geq 0$ such that $n = 2^m$, there exists a symmetric gossip algorithm reaching global finite-time convergence in $O(n \log n)$ steps. Otherwise, for almost all initial conditions, finite-time convergence is impossible via gossip averaging. We also showed that if asymmetric gossiping is allowed, there always exists a globally finite-time convergent gossip algorithm despite the number of nodes. An "all-or-nothing" theorem was established on the finite-time convergence of distributed averaging algorithms indicating that for a class of averaging algorithms defined by sequence selections within a countable set of stochastic matrices, finite-time convergence either holds for all initial conditions, or fails for almost all initial conditions. The results add to the fundamental understanding of distributed averaging algorithms. More challenges lie in the complexity of finding finite-time convergent algorithms in general, and how to construct such an algorithm in a distributed way.

References

[1] G. Latouche, V. Ramaswami. *Introduction to Matrix Analytic Methods in Stochastic Modeling.* 1st edition, ASA SIAM, 1999.

[2] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking, "Randomized rumor spreading," in *Proc. Symp. Foundations of Computer Science*, pp. 564-574, 2000.

[3] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. Conf. Foundations of Computer Science*, pp. 482-491, 2003.

[4] S. Boyd, P. Diaconis and L. Xiao, "Fastest mixing Markov chain on a graph," *SIAM Review*, Vol. 46, No. 4, pp. 667-689, 2004.

[5] S. Boyd, A. Ghosh, B. Prabhakar and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Information Theory*, vol. 52, no. 6, pp. 2508-2530, 2006.

[6] D. Shah, "Gossip Algorithms," *Foundations and Trends in Networking*, vol. 3, no. 1, pp. 1-125, 2008.

[7] D. Mosk-Aoyama and D. Shah, "Fast distributed algorithms for computing separable functions," *IEEE Transactions on Information Theory*, vol.55, no.7, pp. 2997-3007, 2008

[8] F. Fagnani and S. Zampieri, "Asymmetric randomized gossip algorithms for consensus," *IFAC World Congress*, Seoul, pp. 9051-9056, 2008.

[9] J. Liu, S. Mou, A. S. Morse, B. D. O. Anderson, and C. Yu, "Deterministic gossiping," *Proceedings of IEEE*, vol. 99, no. 9, pp. 1505-1524, 2011.

[10] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proceedings of IEEE*, vol. 98, no. 11, pp. 1847-1864, 2010.

[11] S. Kar and J. M. F. Moura, "Convergence rate analysis of distributed gossip (linear parameter) estimation: fundamental limits and tradeoffs," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp.674-690, 2011.

[12] J. Lavaei and R.M. Murray, "Quantized consensus by means of gossip algorithm," *IEEE Trans. Autom. Control*, vol. 57, no.1, pp. 19-32, 2012.

[13] F. Bénézit, A. G. Dimakis, P. Thiran, and M. Vetterli, "Order-optimal consensus through randomized path averaging," *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 5150-5167, 2010.

[14] J. Hajnal, "Weak ergodicity in non-homogeneous markov chains," *Proc. Cambridge Philos. Soc.*, no. 54, pp. 233-246, 1958.

[15] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. 31, pp. 803-812, 1986.

[16] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom.Control*, vol. 48, no. 6, pp. 988-1001, 2003.

[17] V. Blondel, J. M. Hendrickx, A. Olshevsky and J. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," *IEEE Conf. Decision and Control*, pp. 2996-3000, 2005.